Examiner-Initiated Interview Summary	Application No.	Applicant(s)	
	10/750,552	DING ET AL.	
	Examiner	Art Unit	
	J. Derek Rutten	2192	

All Participants:

(1) *J. Derek Rutten*.

(2) *David Edmondson, Reg. No. 35,126*.

Status of Application: *Allowable*

(3) _____.

(4) _____.

Date of Interview: *7 November 2007*

Time: _____

Type of Interview:
☒ Telephonic
☐ Video Conference
☐ Personal (Copy given to: ☐ Applicant ☐ Applicant's representative)

Exhibit Shown or Demonstrated: ☒ Yes ☐ No
If Yes, provide a brief description: *Proposed Claim Amendments.*

Part I.

Rejection(s) discussed:

Claims discussed:
2, 4, 7, 11, 17, 19, 21, 24, 28, 34, 35

Prior art documents discussed:

Part II.

SUBSTANCE OF INTERVIEW DESCRIBING THE GENERAL NATURE OF WHAT WAS DISCUSSED:
See Continuation Sheet

Part III.

☒ It is not necessary for applicant to provide a separate record of the substance of the interview, since the interview directly resulted in the allowance of the application. The examiner will provide a written summary of the substance of the interview in the Notice of Allowability.

☐ It is not necessary for applicant to provide a separate record of the substance of the interview, since the interview did not result in resolution of all issues. A brief summary by the examiner appears in Part II above.

(Examiner/SPE Signature)

(Applicant/Applicant's Representative Signature – if appropriate)

Continuation of Substance of Interview including description of the general nature of what was discussed: Mr. Rutten contacted Mr. Edmondson to suggest an examiner's amendment of independent claims 7, 11, 17, 24, 28, 34, and 35 to include patentable subject matter identified in claims 2 and 4 (likewise claims 19 and 21). Mr. Rutten suggested including limitations from claim 4 into claims 7, 11, 17, 24, 28, 34, and 35. Mr. Edmondson agreed to this amendment but also wanted limitations from claim 2 in the claims. Mr. Edmondson then suggested creating additional new claims reflecting the current claim tree that would feature patentable limitations from claim 2. Mr. Rutten agreed that such claims would be patentable and agreed to include the additional claims.

BLANK ROME LLP
COUNSELORS AT LAW

Watergate • 600 New Hampshire Ave N.W. • Washington, DC 20037
202.772.5800 • fax 202.772.5858

FAX TRANSMITTAL FORM

Please reply to the fax number listed below, if provided.

TO:	J. Rutten	DATE:	08 Nov 2007
COMPANY:	USPTO AU 2192		
FAX:	1-571-273-3703	PAGES:	Cover Plus 19
FROM:	David J. Edmondson	PHONE:	202-772-5838
		FAX:	202-572-1438
RE:	for discussion purposes in 10/750,552		

MESSAGE

Per our discussion this morning. Thank you for your time and consideration.

1. (Canceled)

2. (Previously Presented) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a); and

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error.

3. (Original) The method of claim 2, wherein the search tree is compressed by (i) determining a capacity of each node in accordance with the reuse distance and the bounded relative error and (ii) merging adjacent ones of the nodes in accordance with the capacities of the nodes.

4. (Previously Presented) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a); and

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

maintaining a trace storing the last access times of the last C accesses of the data, where C is a cut-off distance; and

maintaining a search tree storing access times other than the last C accesses, each node in the search tree having a capacity B, where B is a bounded absolute error.

5. (Canceled)

6. (Canceled)

7. (Currently Amended) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

maintaining a trace storing the last access times of the last C accesses of the data, where C is a cut-off distance; and

maintaining a search tree storing access times other than the last C accesses, each node in the search tree having a capacity B, where B is a bounded absolute error; and

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein step (d) comprises forming a reuse distance histogram of the reuse distances by absolute ranges of the reuse distances, and wherein step (d) further comprises forming a reference histogram of the reuse distances by percentile ranges of the reuse distances.

8. (Original) The method of claim 7, wherein the reference histogram is formed for a plurality of training inputs.

9. (Original) The method of claim 8, wherein step (d) further comprises using the reference histograms for the plurality of training inputs to map data size to the reuse distance.

10. (Original) The method of claim 9, wherein the data size is mapped to the reuse distance through linear fitting.

11. (Currently Amended) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

maintaining a trace storing the last access times of the last C accesses of the data, where C is a cut-off distance; and

maintaining a search tree storing access times other than the last C accesses, each node in the search tree having a capacity B, where B is a bounded absolute error;

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein step (d) comprises forming a reuse distance histogram of the reuse distances by absolute ranges of the reuse distances; and

(e) from the reuse distance histogram, forming an affinity group of at least two data which are always accessed within a distance k of one another, wherein k is a predetermined quantity.

12. (Original) The method of claim 11, wherein step (e) comprises selecting the data in the affinity group such that the data in the affinity group have average reuse distances which fulfill a necessary condition with respect to k.

13. (Currently Amended) The method of claim ~~15~~ 12, wherein the necessary condition is that the average reuse distances differ by no more than k.

14. (Original) The method of claim 12, wherein:

the reuse distance histogram comprises B bins; and

the necessary condition is that differences between the average reuse distances, summed over all of the bins, do not exceed kB.

15. (Original) The method of claim 14, wherein step (e) comprises:

(i) initially treating each of the data as an affinity group;

(ii) traversing all of the affinity groups and merging any two affinity groups for which the necessary condition is met; and

(iii) performing step (e)(ii) until no more of the affinity groups can be merged.

16. (Original) The method of claim 11, wherein step (e) is performed a plurality of times for different values of k.

17. (Currently Amended) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

maintaining a trace storing the last access times of the last C accesses of the data, where C is a cut-off distance; and

maintaining a search tree storing access times other than the last C accesses, each node in the search tree having a capacity B, where B is a bounded absolute error;

(d) comparing reuse signatures of the data to determine whether two or more of the data have reuse signatures which differ by less than a predetermined percentage; and

(e) for any two or more of the data whose reuse signatures differ by less than said predetermined percentage, identifying a reference affinity among said two or more data.

18. (Canceled)

19. (Previously Presented) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and

a processor, in communication with the memory, for:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a); and

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein the processor performs step (c) by:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error.

20. (Original) The computing device of claim 19, wherein the search tree is compressed by (i) determining a capacity of each node in accordance with the reuse distance and the bounded relative error and (ii) merging adjacent ones of the nodes in accordance with the capacities of the nodes.

21. (Previously Presented) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and

a processor, in communication with the memory, for:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a); and

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein the processor performs step (c) by:

maintaining a trace storing the last access times of the last *C* accesses of the data, where *C* is a cut-off distance; and

maintaining a search tree storing access times other than the last *C* accesses, each node in the search tree having a capacity *B*, where *B* is a bounded absolute error.

22. (Canceled)

23. (Canceled)

24. (Currently Amended) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and

a processor, in communication with the memory, for:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

maintaining a trace storing the last access times of the last *C* accesses of the data, where *C* is a cut-off distance; and

maintaining a search tree storing access times other than the last *C* accesses, each node in the search tree having a capacity *B*, where *B* is a bounded absolute error; and

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein the processor performs step (d) by forming a reuse distance histogram of the reuse distances by

absolute ranges of the reuse distances, and wherein the processor further performs step (d) by forming a reference histogram of the reuse distances by percentile ranges of the reuse distances.

25. (Original) The computing device of claim 24, wherein the reference histogram is formed for a plurality of training inputs.

26. (Original) The computing device of claim 25, wherein the processor performs step (d) further by using the reference histograms for the plurality of training inputs to map data size to the reuse distance.

27. (Original) The computing device of claim 26, wherein the data size is mapped to the reuse distance through linear fitting.

28. (Currently Amended) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and

a processor, in communication with the memory, for:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

maintaining a trace storing the last access times of the last C accesses of the data, where C is a cut-off distance; and

maintaining a search tree storing access times other than the last C accesses, each node in the search tree having a capacity B, where B is a bounded absolute error;

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein the processor performs step (d) by forming a reuse distance histogram of the reuse distances by absolute ranges of the reuse distances, and

(e) from the reuse distance histogram, forming an affinity group of at least two data which are always accessed within a distance k of one another, wherein k is a predetermined quantity.

29. (Original) The computing device of claim 28, wherein the processor performs step (e) by selecting the data in the affinity group such that the data in the affinity group have average reuse distances which fulfill a necessary condition with respect to k.

30. (Original) The computing device of claim 29, wherein the necessary condition is that the average reuse distances differ by no more than k.

31. (Original) The computing device of claim 29, wherein:

the reuse distance histogram comprises B bins; and

the necessary condition is that differences between the average reuse distances, summed over all of the bins, do not exceed kB.

32. (Original) The computing device of claim 31, wherein the processor performs step (e) by:

(i) initially treating each of the data as an affinity group;

(ii) traversing all of the affinity groups and merging any two affinity groups for which the necessary condition is met; and

(iii) performing step (e)(ii) until no more of the affinity groups can be merged.

33. (Original) The computing device of claim 32, wherein step (e) is performed a plurality of times for different values of k.

34. (Currently Amended) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and

a processor, in communication with the memory, for:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

maintaining a trace storing the last access times of the last C accesses of the data, where C is a cut-off distance; and

maintaining a search tree storing access times other than the last C accesses, each node in the search tree having a capacity B, where B is a bounded absolute error;

(d) comparing reuse signatures of the data to determine whether two or more of the data have reuse signatures which differ by less than a predetermined percentage; and

(e) for any two or more of the data whose reuse signatures differ by less than said predetermined percentage, identifying a reference affinity among said two or more data.

35. (Currently Amended) A method for analyzing affinities among a plurality of events, the method comprising:

(a) monitoring occurrences of the events;

(b) determining a reoccurrence distance for each event, the reoccurrence distance being a number of distinct ones of the plurality of events which occur between two occurrences of said each event, wherein step (b) comprises:

determining a last occurrence of each of the events;

maintaining a trace storing the last occurrences of the last *C* events, where *C* is a cut-off distance; and

maintaining a search tree storing occurrences other than the last *C* occurrences, each node in the search tree having a capacity *B*, where *B* is a bounded absolute error; and

(c) determining, from the reoccurrence distance determined in step (b), an affinity among at least two of the events, the affinity being a tendency of said at least two of the events to occur together.

36. (Original) The method of claim 35, wherein step (c) comprises determining the affinity such that said events always occur within a distance *k* of each other, wherein *k* is a predetermined quantity and the distance is a number of distinct events occurring between occurrences of said at least two of the events.

37. (Original) The method of claim 35, wherein step (c) comprises comparing reoccurrence signatures of the events to determine whether said two or more of the events have reoccurrence signatures which differ by less than a predetermined percentage.

38. (New) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error; and

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein step (d) comprises forming a reuse distance histogram of the reuse distances by absolute ranges of the reuse distances, and wherein step (d) further comprises forming a reference histogram of the reuse distances by percentile ranges of the reuse distances.

39. (New) The method of claim 38, wherein the reference histogram is formed for a plurality of training inputs.

40. (New) The method of claim 39, wherein step (d) further comprises using the reference histograms for the plurality of training inputs to map data size to the reuse distance.

41. (New) The method of claim 40, wherein the data size is mapped to the reuse distance through linear fitting.

42. (New) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error;

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein step (d) comprises forming a reuse distance histogram of the reuse distances by absolute ranges of the reuse distances; and

(e) from the reuse distance histogram, forming an affinity group of at least two data which are always accessed within a distance k of one another, wherein k is a predetermined quantity.

43. (New) The method of claim 42, wherein step (e) comprises selecting the data in the affinity group such that the data in the affinity group have average reuse distances which fulfill a necessary condition with respect to k.

44. (New) The method of claim 43, wherein the necessary condition is that the average reuse distances differ by no more than k.

45. (New) The method of claim 43, wherein:

the reuse distance histogram comprises B bins; and

the necessary condition is that differences between the average reuse distances, summed over all of the bins, do not exceed kB.

46. (New) The method of claim 45, wherein step (e) comprises:

(i) initially treating each of the data as an affinity group;

(ii) traversing all of the affinity groups and merging any two affinity groups for which the necessary condition is met; and

(iii) performing step (e)(ii) until no more of the affinity groups can be merged.

47. (New) The method of claim 42, wherein step (e) is performed a plurality of times for different values of *k*.

48. (New) A method for analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device having a memory in which the data are stored and from which the data are accessed, the method comprising:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error;

(d) comparing reuse signatures of the data to determine whether two or more of the data have reuse signatures which differ by less than a predetermined percentage; and

(e) for any two or more of the data whose reuse signatures differ by less than said predetermined percentage, identifying a reference affinity among said two or more data.

49. (New) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and

a processor, in communication with the memory, for:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error; and

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein the processor performs step (d) by forming a reuse distance histogram of the reuse distances by absolute ranges of the reuse distances, and wherein the processor further performs step (d) by forming a reference histogram of the reuse distances by percentile ranges of the reuse distances.

50. (New) The computing device of claim 49, wherein the reference histogram is formed for a plurality of training inputs.

51. (New) The computing device of claim 50, wherein the processor performs step (d) further by using the reference histograms for the plurality of training inputs to map data size to the reuse distance.

52. (New) The computing device of claim 51, wherein the data size is mapped to the reuse distance through linear fitting.

53. (New) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and

a processor, in communication with the memory, for:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error;

(d) determining a reuse pattern from the reuse distances determined in step (c), wherein the processor performs step (d) by forming a reuse distance histogram of the reuse distances by absolute ranges of the reuse distances, and

(e) from the reuse distance histogram, forming an affinity group of at least two data which are always accessed within a distance k of one another, wherein k is a predetermined quantity.

54. (New) The computing device of claim 53, wherein the processor performs step (e) by selecting the data in the affinity group such that the data in the affinity group have average reuse distances which fulfill a necessary condition with respect to k.

55. (New) The computing device of claim 54, wherein the necessary condition is that the average reuse distances differ by no more than k.

56. (New) The computing device of claim 54, wherein:

the reuse distance histogram comprises B bins; and

the necessary condition is that differences between the average reuse distances, summed over all of the bins, do not exceed kB.

57. (New) The computing device of claim 56, wherein the processor performs step (e) by:

(i) initially treating each of the data as an affinity group;

(ii) traversing all of the affinity groups and merging any two affinity groups for which the necessary condition is met; and

(iii) performing step (e)(ii) until no more of the affinity groups can be merged.

58. (New) The computing device of claim 57, wherein step (e) is performed a plurality of times for different values of k.

59. (New) A computing device capable of analyzing reuse patterns of accesses of data by a program running on a computing device, the computing device comprising:

a memory in which the data are stored and from which the data are accessed; and

a processor, in communication with the memory, for:

(a) running the program on the computing device;

(b) monitoring the accesses of the data by the program during step (a);

(c) determining a reuse distance for each datum from among the data accessed by the program during step (a), the reuse distance being a number of distinct data which are accessed between two accesses of the datum, wherein step (c) comprises:

determining a last access time of each of the data;

organizing a search tree from the last accesses, wherein the search tree comprises a node for each of the data, the node comprising the last access time and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error;

(d) comparing reuse signatures of the data to determine whether two or more of the data have reuse signatures which differ by less than a predetermined percentage; and

(e) for any two or more of the data whose reuse signatures differ by less than said predetermined percentage, identifying a reference affinity among said two or more data.

60. (New) A method for analyzing affinities among a plurality of events, the method comprising:

(a) monitoring occurrences of the events;

(b) determining a reoccurrence distance for each event, the reoccurrence distance being a number of distinct ones of the plurality of events which occur between two occurrences of said each event, wherein step (b) comprises:

determining a last occurrence of each of the events;

organizing a search tree from the last occurrences, wherein the search tree comprises a node for each of the events, the node comprising the last occurrence and a weight of a sub-tree of the node; and

compressing the search tree in accordance with a bounded relative error; and

(c) determining, from the reoccurrence distance determined in step (b), an affinity among at least two of the events, the affinity being a tendency of said at least two of the events to occur together.

61. (New) The method of claim 60, wherein step (c) comprises determining the affinity such that said events always occur within a distance k of each other, wherein k is a predetermined quantity and the distance is a number of distinct events occurring between occurrences of said at least two of the events.

62. (New) The method of claim 60, wherein step (c) comprises comparing reoccurrence signatures of the events to determine whether said two or more of the events have reoccurrence signatures which differ by less than a predetermined percentage.